

# High Fidelity Simulations of Unsteady Flow Through Turbopumps and Flowliners

Cetin C. Kiris, Dochan Kwak, and William Chan  
*NASA Advanced Supercomputing Division (NAS)*  
*NASA Ames Research Center, Moffett Field, CA 94035*

and

Jeffrey A. Housman  
*Graduate Group in Applied Mathematics*  
*U.C. Davis, Davis, CA*

**High fidelity computations were carried out to analyze the orbiter LH2 feedline flowliner. Computations were performed on the *Columbia* platform which is a 10,240-processor supercluster consisting of 20 Altix nodes with 512 processors each. Various computational models were used to characterize the unsteady flow features in the turbopump, including the orbiter Low-Pressure-Fuel-Turbopump (LPFTP) inducer, the orbiter manifold and an experimental test article used to represent the manifold. Unsteady flow originating from the orbiter LPFTP inducer is one of the major contributors to the high frequency cyclic loading that results in high cycle fatigue damage to the gimbal flowliners just upstream of the LPFTP. The flow fields for the orbiter manifold and representative test article are computed and analyzed for similarities and differences. An incompressible Navier-Stokes flow solver INS3D, based on the artificial compressibility method, was used to compute the flow of liquid hydrogen in each test article.**


## I. Introduction

LIQUID rocket turbopumps operate under severe conditions and at very high rotational speeds. The orbiter low-pressure-fuel-turbopump creates transient flow features such as reverse flows, tip clearance effects, secondary flows, vortex shedding, junction flows, and cavitation effects. Flow unsteadiness originating from the orbiter Low-Pressure-Fuel-Turbopump (LPFTP) inducer is one of the major contributors to the high frequency cyclic loading that results in high cycle fatigue damage to the gimbal flowliners just upstream of the LPFTP. The reverse flow generated at the tip of the inducer blades travels upstream and interacts with the bellows cavity. In order to characterize various aspects of the flowfield near the flowliner, various computational models have been developed and high-fidelity computations have been carried out. Results including 14 inducer rotations of a straight pipe model with the LPFTP inducer and 12 inducer rotations of the LPFTP inducer with the addition of upstream and downstream flowliners including 38 slots, an overhang area between the liners, and the bellows cavity are reported in this paper. The incompressible Navier-Stokes flow solver based on the artificial compressibility method was used to compute the flow of liquid hydrogen in each test article. All computations included tip leakage effects with a radial tip clearance of 0.006 inches, a pump operating condition of 104.5% RPL power-level, a mass flow rate of 154.7 lbm/sec, and a rotational speed of 15,761 RPM. The findings include, a significant time-periodic back-flow generated by the inducer reaching 15-20% of the tip velocity and a jet flow of 10-15% of the inducer tip speed which penetrates into the bellows cavity creating an unsteady recirculation region in the cavity. The reverse flow and unsteady recirculation regions create an unsteady interaction between the duct and the bellows cavity resulting in high frequency cycle loading. The back-flow also creates swirl in the bellows cavity on the order of 10% of the inducer tip velocity. The LPFTP inducer with flowliner geometry is further analyzed and 12 inducer rotations have been completed to better represent the time-periodic nature of the flow field.


Additionally, the flow through the orbiter fuel feedline manifold and the flow through the experimental test article are computed. The objective of simulating the flow field through the manifold and the test article is to characterize the similarities and differences between the two configurations.

## II. Computational Models

The first computational model (**Model I**) includes the LPFTP inducer with 4 long and 4 short blades, and a straight duct, which extends 4 duct diameters upstream of the inducer. The bull nose of the inducer, and split blades are included in the model. The objective of studying the inducer model alone is to compare unsteady pressure values against existing data. To resolve the complex geometry in relative motion, an overset grid approach is employed. The geometrically complex body is decomposed into a number of simple grid components. Connectivity between neighboring grids is established by interpolation at each grid outer boundaries. Addition of new components to the system and simulation of arbitrary relative motion between multiple bodies are achieved by establishing new connectivity without disturbing the existing grids. This computational grid has 57 overset zones with 26.1 million grid points.




(a)



(b)

**Figure 1.** (a) Surface grids for LPFTP inducer and the Liquid LH<sub>2</sub> flowliner. (b) Details of the flowliner overset grid system.

The second computational grid system (**Model II**) is based on the first with the addition of the flowliner geometry. The grid system includes 38 upstream slots, 38 downstream slots, the overhang area between liners and the bellows cavity. This model is very similar to the ground test article. It consists of 264 overlapped grids with 65.9 million grid points. Details of the grid system are shown in figures 1(a) and 1(b). The flowliner component consists of an axisymmetric chamber around the external wall of the pipe, and two rows of slots in the streamwise direction. Each slot is a rectangular shaped hole with rounded corners. On the outside wall of the chamber are the bellows, which are shaped like 10 periods of a sine wave. The bellows cavity is connected to the duct via the overhang area and the slots. Two-dimensional overset grids are first created for the bellows, side walls and the overhang area of the bellows cavity. These are then revolved 360 degrees to form the volume grids. Each slot consists of a body-conforming grid and a warped Cartesian core grid in the middle of the hole. The flowliner component alone contains 212 grids and 41 million points.



**Figure 2.** Detail of the computational model for the orbiter fuel feedline manifold.

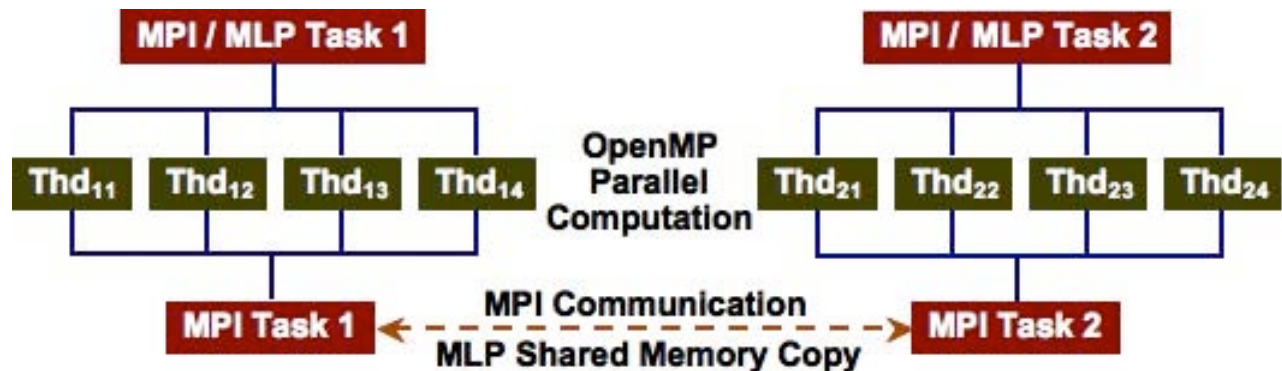
The orbiter fuel feedline manifold grid system consists of an inflow pipe, the manifold, three exit pipes with elbows to the main engines, and two short exit pipes; one for the re-circulating pump and the other for the fill and drain line, (see figure 2). The re-circulating pump and fill and drain line are not included in the computational model so we have closed these two exit pipes. All pipes are connected to the manifold via internal collar grids. The upstream side of the inflow pipe and the downstream side of the exit pipes are modeled by body-conforming O-grids with a singular polar axis running down the core of the pipe. In the regions of the inflow and exit pipes that are near the manifold, the singular axis is avoided by adding a warped Cartesian core grid. Body-conforming grids are used for the walls of the manifold. A series of uniform Cartesian grids are used to occupy the core of the manifold. The entire grid system consists of 38 grids and 12 million points, (see figure 2). A separate computational model was generated for the representative manifold test article. The computational grid representing the test article is a created using O-H grid consisting of six overlapping zones and a total of 7.1 million grid points

In order to speed up and automate the grid generation procedure, script systems has been developed to automatically

and rapidly perform the various steps of the grid generation process prior to the use of the flow solver. Special procedures were developed to automatically create grids for each component type. The component types included in the script are blade, pipe, ring, nose, flowliner, and strut. The blade component is one of the most common parts of a liquid rocket subsystem and may contain multiple sections of one or more sets of different blades, e.g. inducer, impeller, diffuser. The pipe and ring components are used to connect different blade components. Pipes can be straight or curved and are bounded by the shroud. Rings can only be straight and are bounded by both the hub and the shroud. The nose component is a cap that fits at the start or end of the hub. The flowliner is a highly complex part with bellows, and slots. The strut component consists of multiple blades connected to brackets at the shroud end and a central hub at the other end. The strut component was not used in the test article, it was used in the flight configuration, which was not included in this paper.

### III. Numerical Method and Parallel Computations

The present computations are performed utilizing the INS3D computer code, which solves the incompressible Navier-Stokes equations for both steady-state and unsteady flows. The numerical solution of the incompressible Navier-Stokes equations requires special attention in order to satisfy the divergence-free constraint on the velocity field. The incompressible formulation does not explicitly yield the pressure field from an equation of state or the continuity equation. One way to avoid the difficulty of the elliptic nature of the equations is to use the artificial compressibility method, developed by Chorin<sup>1</sup>. The artificial compressibility method introduces a pseudo-time derivative of the pressure term into the continuity equation. This transforms the elliptic-parabolic type partial differential equations into the hyperbolic-parabolic type. An incompressible flow solver, INS3D, has been developed<sup>2-4</sup> based on this formulation. Since the convective terms of the resulting equations are hyperbolic, upwind differencing can be applied to these terms. The current version uses Roe's flux-difference splitting<sup>5</sup>. The third and fifth-order upwind biased differencing used here are implementations of a class of high-accuracy flux-differencing schemes for the compressible flow equations. To obtain time-accurate solutions, the equations are iterated to convergence in pseudo-time for each physical time step until the divergence of the velocity field has been reduced below a specified tolerance value. The total number of sub-iterations required varies depending on the problem, time step size and the artificial compressibility parameter used. Typically the number ranges from 10 to 30 sub-iterations. The linearized system of equations are solved iteratively by using a non-factored symmetric Gauss-Seidel type line-relaxation scheme<sup>6</sup>, which is linearly stable and allows a large pseudo-time step to be taken. Computer memory requirements for the flow solver INS3D with line-relaxation is 45 times the number of grid points in words.




**Figure 3.** MLP and MPI hybrid parallel organization for INS3D.

Two distinct parallel processing paradigms have been implemented into the INS3D code. These include the Multi-Level Parallel (MLP)<sup>7</sup> and the MPI/OpenMP hybrid parallel programming models. Both models contain coarse and fine grain parallelism. Coarse grain parallelism is achieved through a UNIX fork in MLP and through explicit message passing in the MPI/OpenMP hybrid code. Fine grain parallelism is achieved using OpenMP<sup>8</sup> compiler




directives in both the MLP and MPI/OpenMP hybrid codes. The multi-level parallel organization for INS3D is shown in figure 3. Both MLP and MPI/OpenMP codes use a group based data structure for global solution arrays. Computations on a single node of the Columbia platform have been carried out using the MLP and MPI/OpenMP hybrid codes. In the MLP implementation, all data communication at the coarsest and finest parallelization levels is accomplished via direct memory referencing instructions. The coarsest level parallelism is supplied by spawning off independent processes via the standard UNIX fork. A library of routines is used to initiate forks, to establish shared memory arenas, and to provide synchronization primitives. The MLP code uses a global shared memory data structure for overset connectivity arrays, while the MPI/OpenMP code uses local copies of the connectivity arrays providing a more local data structure. Computations were performed to compare the scalability between the MLP and MPI/OpenMP hybrid versions of the INS3D code on the Columbia system using the BX2b processors. Initial computations using one group and one thread were used to establish the baseline runtime for one physical time step, where 720 such time steps are required to complete one inducer rotation. Figure 4 displays both the time per iteration (in minutes) versus the number of CPUs as well as the speedup factor for both the MLP and MPI/OpenMP hybrid versions of the code. Here 36 groups have been chosen to maintain good load balance for both versions. Then the runtime per physical time step is obtained using various numbers of OpenMP threads (1, 2, 4, 8, and 14).



**Figure 4.** INS3D-MLP and MPI+OMP performance on the Columbia platform.


The scalability for a fixed number of both MLP and MPI groups and varying OpenMP threads is good, but begins to decay as the number of OpenMP threads becomes large. Further scaling can be accomplished by fixing the number of OpenMP threads and increasing the number of MLP/MPI groups until the load balancing begins to fail. Unlike varying the OpenMP threads that does not affect the convergence rate of INS3D, varying the number of groups may deteriorate the convergence rate. This will lead to more iterations even though faster runtime per iteration is achieved. Figure 4 shows that the MLP and MPI/OpenMP codes perform almost equivalently for one OpenMP thread, then as the number of threads are increased the MPI/OpenMP hybrid version of the code begins to perform slightly better than the MLP version. This advantage can be attributed to having local copies of the connectivity arrays in the MPI/OpenMP hybrid code. Having the MPI/OpenMP version of INS3D as scalable as the MLP version is promising since this version is easily portable to other platforms.

Performance results of the INS3D MPI/OpenMP code on multiple BX2b nodes are compared against single node results. The results include running the MPI/OpenMP version using two different communication paradigms, master-worker communication and point-to-point communication. The runtime per physical time step is recorded using 36 MPI groups and 1, 4, 8 and 14 OpenMP threads on one, two, and four BX2b nodes. The communication between nodes is achieved using the InfiniBand and NUMalink4 interconnects denoted IB and XPM respectively. Figure 5 contains the results using the point-to-point communication paradigm. When comparing the performance of using multiple nodes with that of the single node we observe that the scalability of the two node and four node runs with NUMalink4 interconnects is similar to the single node runs which also use NUMalink4. When using InfiniBand interconnects we observe a 10-29 % increase in runtime per iteration on two and four node runs. The difference in runtime per iteration between four node runs and two node runs decreases as the number of CPUs increases. Figure 6 displays the results using the master-worker communication paradigm. The first observation is that the time per iteration is much higher using this communication protocol compared to the point-to-point communication. We also see a more significant deterioration of the scalability between the single and multiple node runs, even for the two and four node runs with NUMalink4 interconnects where almost no difference is observed when using point-to-point communication.



**Figure 5.** Performance of INS3D across multiple BX2b nodes using NUMalink4 and InfiniBand interconnection (MPI point-to-point communication).

Using NUMalink4 interconnects we observe a 5-10% increase in runtime per iteration for one to two nodes and an 8-16% increase using four nodes. This is because the master resides on one node and all workers on the other nodes must communicate with the master using the interconnect. Alternatively, when using point-to-point communication many of the messages are bound to the node from which they are sent. In fact the MPI groups can be manipulated so that a minimum (as low as one in many cases) number of messages must be passed between each node. Note, this optimization has not been utilized here and will be studied further. An additional 14-27% increase in runtime is observed when using InfiniBand interconnects instead of NUMalink4, a similar increase was observed when using point-to-point communication, which is shown in figure 5.




**Figure 6.** Performance of INS3D across multiple BX2b nodes using NUMalink4 and InfiniBand interconnection (MPI master-worker communication).


#### IV. Computed Results

In all computations, the inducer tip leakage effect is included with a tip clearance of 0.006 inches. The pump operating condition is 104.5% RPL power-level with a mass flow of 154.7 lbm/sec and a rotational speed of 15,761 RPM. The problem was non-dimensionalized with a reference length of one inch and reference velocity equal to the inducer tip speed. The Reynolds number for these calculations is 36 Million. Liquid Hydrogen is treated as an incompressible single-phase fluid. The past decade has seen considerable progress in the development of engineering CFD models for the multiphase flows characteristic of cavitation. The most practical approach amongst these is the homogeneous-mixture model, wherein the liquid-vapor mixture is treated using individual transport equations for each phase and appropriate source terms are employed to describe the phase-change process. In spite of the progress made in multi-phase simulations in recent years, cavitation remains an extremely complex physical phenomenon. It is therefore crucial to validate these models further and obtain a confidence level before applying them to complex engineering problems.

An initial condition of flow at rest and no inducer rotation is used to start the computations. Then, the inducer is rotated at full speed as a start-up procedure. Mass flow is specified at the inflow and characteristic boundary conditions are used at the outflow. Simulations for 14 inducer rotations were completed for model I, and 12 inducer rotations were completed for model II. The time history of non-dimensional pressure difference from INS3D calculations (model I) at a location where experimental measurements are taken is plotted in figure 7(a). Even though computed results have not fully converged to periodicity in time and may still show the evidence of start-up transients, the dominant 4N unsteadiness at a fixed location is shown in figure 7(a). Figure 7(b) shows maximum and minimum non-dimensional pressure values recorded from the experimental data. Comparisons between CFD results and hot fire test data also show good correlation in the non-dimensional pressure amplitudes.




**Figure 7.** Time history of non-dimensional pressure during one inducer rotation (model I, 14<sup>th</sup> inducer rotation), and Min/Max values of non-dimensional pressure in hot fire test.




**Figure 8.** Instantaneous surface pressure contours on inducer and flowliner (Model II).



Figure 8 shows the instantaneous non-dimensional pressure contours on the inducer and the flowliner surfaces (Model II). Blue color indicates the least and magenta indicates the greatest value. The pressure difference between the pressure side (facing downstream into the pump) and suction side (facing upstream) of the inducer blade is clearly seen in these pictures. The back flow near the inducer blade tip is caused by this pressure difference.




**Figure 9.** Axial velocity contours at an instantaneous time in a vertical cut plane.




**Figure 10.** Velocity vectors colored by total velocity magnitude.

The region of reverse flow extends far enough upstream to interfere with both flowliners in the gimbal joint. Figure 9 shows the extend of the backflow at an instantaneous time. Positive axial velocity values are colored in red and negative axial velocity values are colored in blue. Axial velocity values are non-dimensionalized by tip velocity. It is clearly seen that there is a strong interference between the backflow in the duct and the flow in the bellows cavity. The backflow velocity magnitude reaches 15-20 % of the tip velocity magnitude in the overhang area between the liners. It should be noted that this interaction is unsteady and backflow travels in the circumferential direction as well. Due to strong interactions in the overhang area, flow is excited in the bellows area that results in time-dependent recirculation regions. This observation can be seen in figure 10, where velocity vectors are plotted in the region near the bellows cavity and overhang area. As shown in figure 10,


strong jet flow, with velocities of about 10-15% of the inducer tip speed, penetrates directly into the bellows cavity resulting in strong unsteady recirculation regions in the cavity. The time-dependent interaction between the duct and the bellows cavity can be one of the major contributors for high cycle loading. Figure 10 also shows that modeling the gap in the overhang area between flowliners is very important. Jet-like flow in the overhang area pushes the fluid in the bellows cavity toward the duct through the slots. Without proper modeling of this detailed geometry, one cannot obtain fine scale flow unsteadiness on the liner. This transient phenomenon creates unsteady pressure loading spectrum on the flowliner surfaces. Backflow also causes pre-swirl to occur in the flow approaching the inducer.




(a)



(b)



(c)

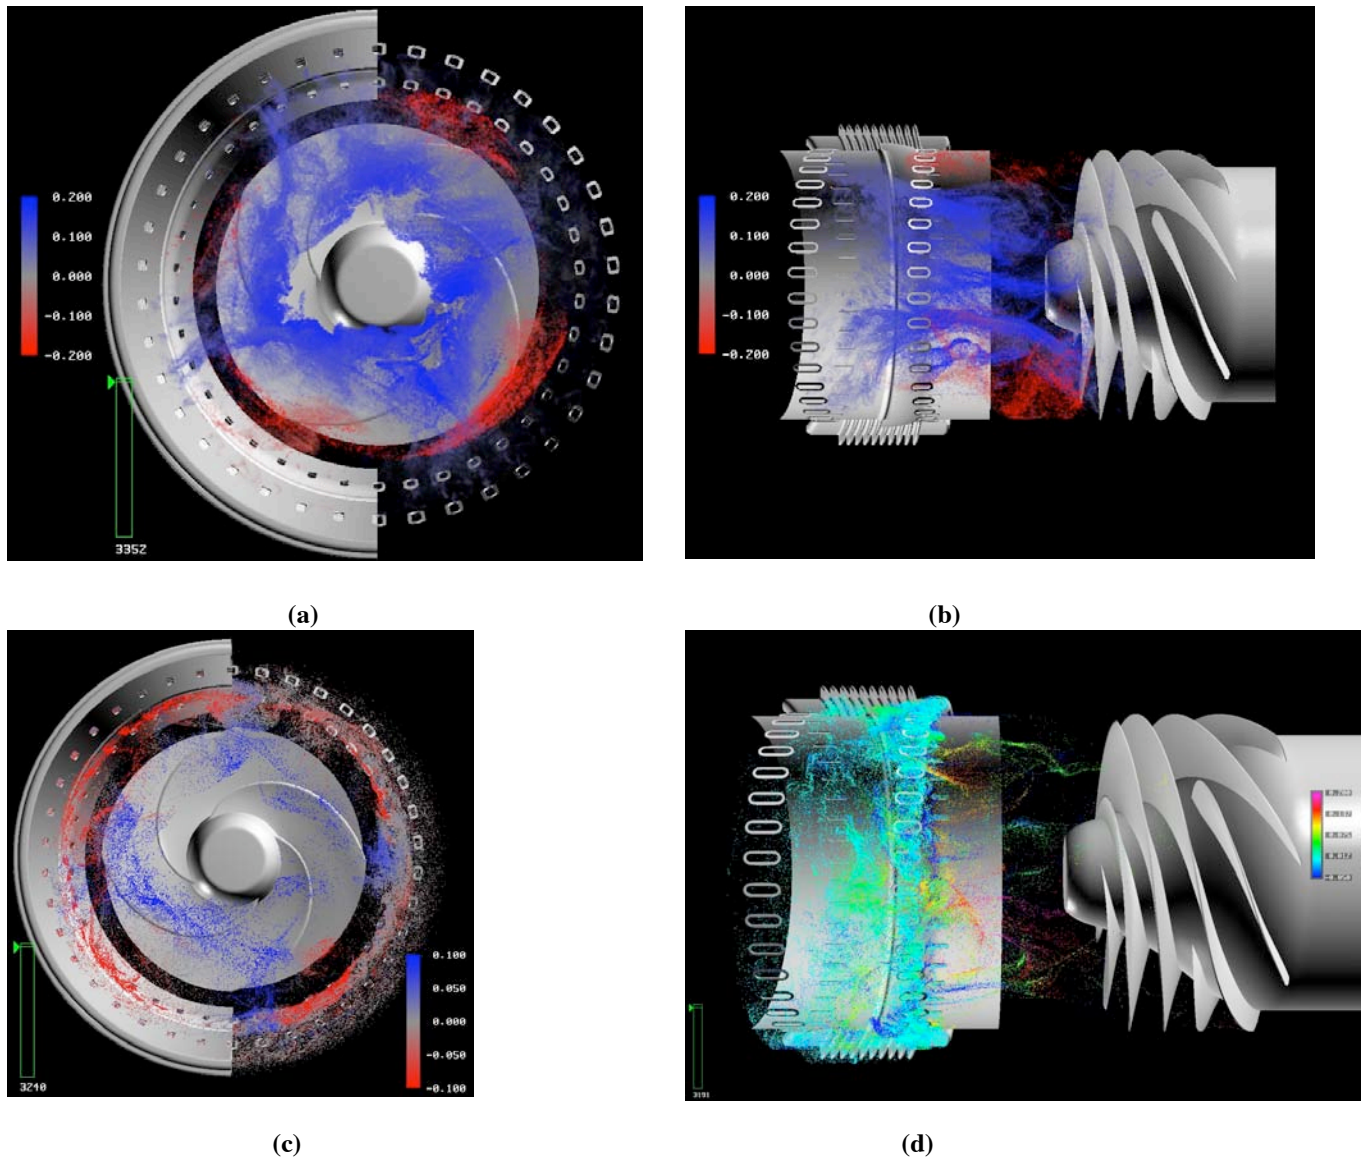


(d)

**Figure 11.** Instantaneous snapshots of particle traces colored by axial velocity values.



In figure 11(a) particles are released from one upstream slot, one downstream slot, and from the overhang area. In figures 11(b)-11(d) the particles are released from each of the 38 downstream slots. The mean flow direction is from left to right in the figure 11(a) and figure 11(c), towards the inducer. In figure 11(b) and figure 11(d) we are viewing the flowliner from the inducer. The figures are snapshots from particle trace animations that have evolved for one and a half rotations. The particles are colored by the axial velocity values. The particles released from both upstream and downstream slots in figure 11(a) are interacting with neighboring slots traveling from the duct to the bellows cavity and from the bellows cavity to the duct. This indicates a highly unsteady interaction between the flow field in the bellows cavity and the flow field in the duct. The swirl present in the particle trace appears to be higher in the downstream slots where particles have traced four and half slots while the upstream particles have traveled only two and half slots. The particles released from the overhang area tend to travel into the bellows cavity with the exception of a few particles traveling into the duct. In figure 11(b) we are observing the particles from the inducers perspective where we can see the particles moving into and out of the bellows cavity through the slots. There are four regions of positive velocity particles traveling from the walls toward the center. These are particles released from the slots that do not interact with the four reverse flow regions created by the inducer. Analyzing the interior of the duct we can see the existence of the swirl near the downstream flowliner.




**Figure 12.** Instantaneous snapshots of particle traces colored by axial velocity values.

Examining figure 11(c) we can see that some of the particles released from the slots travel to the overhang region where they interact with the reverse flow and some of the particles become trapped in the bellows cavity. This suggests that some particles are recirculating through the overhang area into the bellows cavity and through the slots into the duct. This unsteady interaction between the four backflow regions generated by the inducer and the secondary flow features in the bellows cavity lead to a break up of the 4-theta forcing function in the flowliner area. Figure 11(d) suggests the existence of a 5-theta forcing function, but more post-processing and future investigations need to be conducted before any definitive conclusion is possible.

In figures 12(a) and 12(b) particles are released from the upstream slots and evolve for five inducer rotations. The color of the particles represent forward flow (blue) and backward flow (red). We see from the figures that the particles are driven towards the center of the duct and travel to the inducer where some of the particles are trapped into the backflow regions with the 4theta forcing function. In figures 12(c) and 12(d) the particles are released from the downstream slots and evolve for five rotations. In figure 12(c) the particles are colored the same as in the previous two figures. We observe a much more complicated flow structure where many of the particles travel into and out of the bellows cavity. In figure 12(d) we have multi-colored the particles with axial velocity in order to illustrate their presence in the bellows cavity. When the particles are released from the downstream slots they are under the influence of the backflow and swirl such that fewer particles travel towards the inducer.


Figure 13 displays the axial velocity contours on a vertical plane. Due to the number of slots, which is 38 instead of 36, there is a lack of symmetry on the vertical plane. To create this figure the data on each of the structured overset grids are cut vertically and projected onto an unstructured two-dimensional surface. Inherent in this process is the creation of small discontinuities in the contours between overset grids that do not line up with one to one matching in the selected two-dimensional plane. The contours show strong reverse flow regions coming from the blade, traveling through the overhang region creating a jet like flow on the order of 10% of the inducer tip speed. The backflow regions travel up from the inducer blades to the upstream flowliner.




**Figure 13.** Instantaneous axial velocity contours on a vertical plane (12<sup>th</sup> rotation).



The flow through the orbiter fuel feedline manifold and the flow through the experimental test article are computed in order to characterize the similarities and differences between the two configurations. Initially, steady-state calculations were conducted for both the orbiter manifold and the representative test article. The calculations for the orbiter manifold did not converge to a steady solution because of high grid resolution which captures the fine scale unsteady details that exist in high Reynolds number flows. Instead, time-accurate calculations were performed and the mean flow results are presented here. For both of these computations the same non-dimensionalization is used as in the flowliner analysis including a Reynolds number of 36 million based on the inducer tip speed and the reference length of an inch. Consistent inflow and exit boundary conditions are used for the test article and the orbiter manifold, such that the test article exit and the three engines downstream of the orbiter manifold receive the same mass-flow rate. At the inflow, the mass-flow rate is specified with a corresponding turbulent velocity profile and the pressure is calculated through the characteristic relation. At the exit the pressure is extrapolated and the mass-flow rate is enforced.



**Figure 14.** Total velocity magnitude contours at the outlet of the orbiter manifold and inlet to the LPFTP.



**Figure 15.** Total velocity magnitude contours at various cross-sections of the test article model.

The orbiter manifold displayed in figure 14 is qualitatively different for each of the three outlet sections, and none of these sections are represented well by the test article results. Figure 14 also shows that a large velocity profile has developed at the outer wall of the engine 1 feedline, while the high velocity region is better distributed around the entire wall of the engine 3 feedline, and high-low velocity regions near the walls of the engine 2 feedline are not well-developed. One of the primary reasons for the different velocity profiles at the three engine feedlines is the difference in lengths of the engine feedlines. Additionally the velocity profiles for each of the three engines are different as they leave the manifold. In figure 15, the velocity magnitude is displayed at various cross-sections of the test article. The flow direction is from the top of the figure to the bottom. The velocity profile is uniform before the first turn. Then it becomes non-uniform after the first turn, but it returns to a more uniform distribution toward the end of the pipe.

## V. Summary

High fidelity turbopump and flowliner calculations were performed on the Columbia supercluster by using both MLP and MPI/OpenMP hybrid implementations of the INS3D code. Single node computations indicated that the MPI/OpenMP version of the code showed similar performance characteristics as the MLP version. Multiple node computations showed that the point-to-point implementation of the MPI/OpenMP code performs more efficiently than the master-worker version of the MPI code. CFD results confirmed that there is a backflow caused by the LPFTP inducer. The region of reverse flow extends far enough upstream to interfere with both flowliners in the gimbal joint. Computed results for the straight duct hot fire test configuration have been verified by correlation with unsteady pressure measurements. The details of the flowfield at instantaneous times are illustrated. CFD results confirmed that there is a strong unsteady interaction between the backflow regions caused by the LPFTP inducer and secondary flow regions in the bellows cavity through the flowliner slots. It is observed that the swirl on the duct side of the downstream flowliner is stronger than on the duct side of the upstream flowliner. Due to this swirl, there are more significant unsteady flow interactions through the downstream slots than those observed in the upstream slots. Inducer rotations are simulated in order to understand the root cause of the flowliner crack problem. The flow through the test article and the flow through the orbiter fuel feedline manifold were simulated without the LPFTP. It was observed that incoming flow to the flowliner and inducer was more uniform in the test article than in the orbiter manifold. Additionally, each engine LPFTP in the orbiter receives significantly different velocity distributions. Because of these differences observed in the computed results it is not possible for the test article to represent the three different engine feedlines simultaneously.

## Acknowledgments

The authors are grateful to Alex Te for his help in CAD work, Tim Sandstrom and David Ellsworth for their help in the visualization work, and to NASA-NESC Flowliner ITA team members for helpful discussions. The authors would like to thank Mr. Fred Martin and Dr. Goetz Klopfer for reviewing this article.

## References

1. Chorin, A., J., "A Numerical Method for Solving Incompressible Viscous Flow Problems," Journal of Computational Physics, Vol. 2, pp. 12-26, 1967.
2. Kiris, C., and Kwak, D., 'Parallel Unsteady Turbopump Simulations for Reusable Launch vehicle,' Frontiers of Computational Fluid Dynamics 2002, Caughey, D.A. and Hafez, M., ed, World Scientific, 2002.
3. Kiris, C., Kwak, D., and Rogers, S., 'Incompressible Navier-Stokes Solvers in Primitive Variables and Their Applications to Steady and Unsteady Flow Simulations,' Numerical Simulations of Incompressible Flows, Hafez, M., ed, World Scientific, 2003.
4. Rogers, S. E., Kwak, D. and Kiris, C., "Numerical Solution of the Incompressible Navier-Stokes Equations for Steady and Time-Dependent Problems," AIAA Journal, Vol. 29, No. 4, pp. 603-610, 1991.
5. Roe, P.L., "Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes," J. of Comp. Phys., Vol. 43, pp. 357-372 1981.
6. McCormack, R., W., "Current Status of Numerical Solutions of the Navier-Stokes Equations," AIAA Paper No. 85-0032, 1985.
7. Taft, J. R., "Achieving 60 Gflop/s on the Production CFD Code OVERFLOW-MLP," Parallel Computing, 27(4):521-536, 2001.
8. H. Jin, M. Frumkin and J. Yan, "Automatic Generation of OpenMP Directives and Its Application to Computational Fluid Dynamics Codes," in the Proceeding of the Third International Symposium on High Performance Computing, Tokyo, Japan, Oct. 16-18, 2000.